

A Case Study of Verifying and Validating an Astrophysical Simulation Code

A. C. Calder^{1,2}, B. Fryxell^{1,3}, T. Plewa^{1,2,4}, R. Rosner^{1,2,3}, L. J. Dursi^{1,2}, V. G. Weirs^{1,2},
T. Dupont^{1,5}, H. F. Robey⁶, J. O. Kane⁶, B. A. Remington⁶, R. P. Drake⁷, G. Dimonte⁶,
M. Zingale^{1,8}, A. Siegel^{1,5}, A. Caceres^{1,9}, K. Riley¹, N. Vladimirova^{1,2},
P. Ricker^{1,10}, F. X. Timmes^{1,2}, K. Olson^{1,11}, and H. M. Tufo^{1,5}

ABSTRACT

We present a case study of verifying and validating FLASH, a parallel, adaptive-mesh hydrodynamics code for studying the compressible, reactive flows found in many astrophysical environments. We briefly describe the astrophysics problems of interest and physics relevant to these problems. We discuss our methodology for verification and validation, and we describe our automated testing process for software verification. We present the result of one verification test that typifies our methodology for algorithmic verification. We present the results of two fluid dynamics validation tests in which we compared simulations to laboratory results. The first is of a laser-driven shock propagating through a multi-layer target, a configuration subject to both Rayleigh-Taylor and Richtmyer-Meshkov instabilities. The second test is a classic Rayleigh-Taylor instability, where a heavy fluid is supported against the force of gravity by a light fluid. Our simulations of the multi-layer target experiments showed good

¹Center for Astrophysical Thermonuclear Flashes, The University of Chicago, Chicago, IL 60637

²Department of Astronomy & Astrophysics, The University of Chicago, Chicago, IL 60637

³Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637

⁴Nicolaus Copernicus Astronomical Center, Bartycka 18, 00716 Warsaw, Poland

⁵Department of Computer Science, The University of Chicago, Chicago, IL 60637

⁶Lawrence Livermore National Laboratory, Livermore, CA 94550

⁷Department of Atmospheric, Oceanic, and Space Sciences, University of Michigan Ann Arbor, MI 48109

⁸Dept of Astronomy and Astrophysics, The University of California, Santa Cruz, Santa Cruz, CA 95064

⁹Department of Physics The University of Chicago, Chicago, IL 60637

¹⁰Department of Astronomy The University of Illinois, Urbana, IL 61801

¹¹UMBC/GEST Center, NASA/GSFC, Greenbelt, MD 20771

agreement with the experimental results, but our simulations of the Rayleigh-Taylor instability did not agree well with the experimental results. We discuss our findings and potential reasons for the discrepancies.

1. Introduction

Astrophysical events are the most violent and energetic events occurring in the Universe. The complexity of most of these problems requires the development of simulation technology for progress in theoretical research. FLASH, a parallel, adaptive-mesh simulation code for the compressible, reactive flows found in many astrophysical environments is one such tool under development. The motivation for the development of FLASH is to advance the solution of several astrophysical problems related to thermonuclear flashes on the surfaces and in the interiors of compact objects. In particular, the problems of interest are type I X-ray bursts, classical novae, and Type Ia supernovae. These events all involve the accretion of material from a companion star onto the surface of the compact star, followed by the ignition of either the core of the compact star or the material accreted onto the surface.

The global physical phenomena common to all three of these events include an accretion flow onto the surfaces of compact stars, shear flow and Rayleigh-Taylor instabilities (Taylor 1950; Chandrasekhar 1981) on the stellar surfaces and in the core, ignition of thermonuclear burning in degenerate matter, development of convection, propagation of nuclear burning fronts, and expansion of the stellar envelope. An understanding of these global phenomena requires knowledge of the fundamental physical processes involved in each. Accordingly, much of our research effort focuses on the basic “microphysics.” These fundamental processes include turbulence at large Reynolds and Rayleigh numbers, fluid instabilities and mixing, convection and the convective penetration of stable matter at high densities, thermodynamics in relativistic and degenerate regimes, the propagation of both subsonic and supersonic burning fronts, self-gravity, and radiation hydrodynamics. The goal is to understand these processes and incorporate this understanding into astrophysical models.

Because most astrophysical “experiments” are limited to observations of distant events, simulating and modeling these phenomena are largely prediction, foretelling the state of a system for which the simulations or model have not been validated (AIAA 1998). Verification and validation are crucial, however, precisely because so much of astrophysical modeling is prediction. For credible predictions, astrophysical simulation codes must demonstrate the ability to meaningfully describe Nature where the codes can be tested.

Progress is made in verifying and validating an astrophysical simulation code by applying

accepted V&V techniques, adjusting the techniques to the particular constraints posed by the astrophysics. Verifying is carefully studying and quantifying the error in simulations of problems for which an exact or accepted solution is known. The process of validation is to identify the key physical processes involved in the events in question, and for each, to carefully compare simulations to the results of experiments probing these processes. The results and experience gained from such closely related cases serves to better calibrate the simulation code and build confidence in subsequent astrophysical predictions. This process of validation for astrophysical codes relies on the fact that the equations that describe the physics scale from microscopic to stellar scales (Ryutov, et al. 1999).

In this manuscript, we present the some of the results of our efforts at verifying and validating the hydrodynamics module in FLASH. We briefly describe the astrophysical problems of interest and the importance of fluid instabilities in these problems. We present the results of one verification test that typifies our approach to verification, the advection of an isentropic vortex. We describe our automated testing process for software verification. We present the results of validating the code with the results from two laboratory experiments designed to probe fluid instabilities that are expected to be found in the astrophysical problems of interest. Complete details of our verification and validation tests may be found in Calder et al. (2002).

1.1. Astrophysical Thermonuclear Flashes

Thermonuclear flashes, events of rapid or explosive thermonuclear burning, occur in a variety of stellar settings. These events include type I X-ray bursts, classical novae, and Type Ia supernovae, all of which involve a close binary system in which matter from a companion star accretes onto the surface of a compact star (neutron star or white dwarf). Either the core of the compact object or the accreted layer on the surface of the compact object ignites under electron-degenerate conditions, and a thermonuclear burning front is born and begins to propagate.

These events provide fantastic observational displays and are also tools with which to potentially answer several fundamental questions. The light curves and spectra of X-ray bursts can provide information about the masses and radii of neutron stars (Lewin, van Paradijs, & Taam 1993; Lamb 2000) and thus also provide information about the nuclear equation of state. Classical novae can provide information about the abundances (mass fractions) of intermediate-mass elements in the universe and the dynamics of white dwarfs in close binary systems (Gehrz et al. 1998). Type Ia supernovae provide additional information about the abundances of intermediate-mass and heavy elements and play a crucial role as

“standard candles” in determining cosmological parameters such as the Hubble constant, H_0 , the mass density, Ω_M , and the cosmological constant or vacuum energy density, Ω_Λ (see Turner 2001, and references therein).

Type I X-ray bursts are flashes that start at the bottom of a very thin layer ($\sim 10 - 100$ m) of hydrogen-rich or helium-rich fuel that has accreted onto the surface of a neutron star (Taam 1985; Lewin, van Paradijs, & Taam 1993; Taam et al. 1993). The total energy released by burning the fuel into ash is a factor of $\sim 20 - 100$ less than the gravitational binding energy so the accreted material is tightly bound to the neutron star and the flash is not quenched by expansion of the envelope. The fuel in the accreted envelope is incinerated all the way to iron-peak or heavier nuclei (Shatz et al. 2001).

Novae result from the ignition of a layer ($\sim 10^4$ m) of hydrogen-rich material that has accreted onto the surface of a white dwarf (Truran 1982; Shara 1989; Starrfield 1989; Livio 1994). The total energy released by thermonuclear burning is a factor of ~ 100 more than the gravitational binding energy so there is an enormous expansion of the white dwarf’s envelope, which engulfs the companion star and forms a common envelope binary.

Type Ia supernovae are thought to be due to carbon flashes that ignite in the cores of accreting white dwarfs (Woosley & Weaver 1986; Nomoto, Yamaoka, & Shiegeyama 1994; Niemeyer 1995; Niemeyer & Hillebrandt 1995). The mechanism of the explosion is not well understood and models include subsonic burning fronts (deflagrations), supersonic burning fronts (detonations), and transitions from deflagrations to detonations (Hillebrandt & Niemeyer 2000, and references therein).

1.2. The Nature of Scientific Simulation

Our ultimate goal is to understand these thermonuclear flashes in as much detail as possible. We use numerical simulation as a complementary approach to observations, which, as explained above, yield limited information. We consider proposed mechanisms for the flashes and use simulations to test these mechanisms by comparing the simulation results to available observational data. In this sense, the simulations are experiments to determine if the mechanisms are valid.

If the observations and the results of simulations agree over an increasing number of comparisons covering an expanding parameter range, our confidence in the proposed mechanism grows. When observations and simulation results differ, we determine and, if possible, address the source of the discrepancies. Sometimes we learn that our proposed mechanism is wrong, and another mechanism must be developed and tested. By this

process, we will eventually understand astrophysical thermonuclear flashes – that is, by using simulation results to validate our hypotheses. For this process to work, we must have confidence in our simulation results.

Another subject worthy of mention is that of the role of models and submodels. Complete simulations of thermonuclear flashes with detailed models, i.e. models starting from first principles (Oran & Boris 2001), is well beyond the reach of available computing power. Accordingly, full simulations of flashes require phenomenological models for many physical processes. For example, resolving a nuclear flame with a detailed model requires ~ 20 grid points spanning < 0.1 millimeter of physical space, while the radius of a white dwarf is ~ 1000 kilometers (Dursi, et al. 2002). Simulating the full star will require a phenomenological or empirical model for the flame.

These submodels must be developed, verified, and validated, just as our proposed mechanisms for flashes. FLASH is used in this capacity as well – simulation results using detailed and phenomenological models for simpler problems are compared with each other, and if possible to observational and physical experimental data. In this paper we describe our efforts at validation for a model for fluid dynamics by comparing simulation results to experimental data.

Thus the nature of scientific simulation is twofold. At some times the simulation results are used as experimental data to test models (theories). At other times, the simulation is the model to be tested by observational, experimental, or other computational data. In both situations, V&V are the processes by which we progress in our understanding.

1.3. The Role of Fluid Instabilities in Flash Problems

Fluid instabilities and mixing play an important role in thermonuclear flash events. In the nova case, determining whether or not there is substantial mixing between the accreted hydrogen-helium envelope and the carbon-oxygen surface layer of the white dwarf is crucial because without the intermediate-mass elements from the white dwarf, hydrogen burning would be too slow to produce a nova. Furthermore, without this mixing it is difficult to produce the observed abundances of intermediate-mass nuclei in the ejecta (see Rosner et al. 2001, and references therein). In Type Ia supernova, a subsonic burning front that begins near the center of a massive white dwarf is subject to Kelvin-Helmholtz, Landau-Darrieus, and Rayleigh-Taylor instabilities (Khokhlov, Oran, & Wheeler 1997; Khokhlov 2001; Hillebrandt & Niemeyer 2000). Growth of these instabilities dramatically increases the surface area of the burning front. This increase in surface area increases both the burning

rate and the speed of the front.

Because fluid instabilities play a fundamental role in thermonuclear flash events and may be probed in reasonably good terrestrial experiments, experiments involving fluid instabilities have been the focus of our validation efforts thus far. In later sections, we present the results of two validation problems involving fluid instabilities—a laser-driven shock propagating through a multi-layer target and a classic Rayleigh-Taylor problem.

1.4. The Process of Verification and Validation

V&V are maturing areas of study, and there is considerable literature available (cf. AIAA 1998; Oberkampf 1998; Pilch et al. 2001; Roache 1998a,b). The fundamental goal of V&V is the assessment of error and uncertainty in a computational simulation. The requisite methodology is complex because it must address sources of error in theory, experiment, and computation. Because these areas of study present diverse perspectives (and because V&V is still a developing field), it is common to find disagreement in the terminology of V&V (AIAA 1998).

1.4.1. Algorithmic Verification

Verification is the process of determining that the implementation of a model accurately represents the developer’s conceptual description of the model (AIAA 1998). Verification begins by identifying the purported design goal of the code or code modules. This is the step Roache refers to as defining “precisely what partial differential equations are being solved,” though the process may require a larger scope. In the language of the AIAA, this may be thought of as identifying and understanding the implementation of the model. Verification, confirming that simulations produced by the code accurately implement the design goal of the code (i.e. the implementation accurately represents the model and the solution to the model), requires identification and a quantitative description of the error. This can be approached in several distinct ways: comparison of results with analytical solutions or highly accurate benchmarks; testing for some form of iterative convergence; “consistency” checks of well understood global quantities (e.g. conservation); mesh and time step convergence studies, as is appropriate for finite volume methods (AIAA 1998). Our verification tests rely mainly on the latter strategy, though we do offer some important comments on our procedure of analytical testing and software quality assurance (“re-verification”) (Roache 1998a).

Our verification methodology is as follows: we perform several simulations of a given

test problem at increasing resolutions, compare the results to the accepted results, and confirm that the solutions are converging to the accepted result at the correct rate. We note that an added level of complexity arises in the case of discontinuous flows because not all numerical methods for solving a given set of PDEs are equivalent. Various methods handle flow discontinuities differently and/or rely on different (intrinsic numerical or physically-motivated) subgrid models. Therefore, confirming that a numerical method is correctly solving a set of PDEs can also be a validation problem.

1.4.2. Some Comments on Software Verification

An important part of the verification process is code testing – what Roache refers to as re-verification (Roache 1998a) and is sometimes referred to as software verification. Assuming the solution is correct at one time, how can we easily reaffirm its correctness at some future time? Simple coding changes are made, the software is ported to new platforms, compilers and external libraries are updated – when all of this occurs simultaneously with more significant development, sources of error can be easily confused. In many ways, this is a far more concrete process than our other verification analyses, yet it is far from trivial. In our experience, such re-verification is a crucial facet of group scientific software development.

The backbone of the FLASH development process is an abstract, extensible testing framework called FLASH_test. FLASH_test defines three basic sets of tests: compilation, comparison, and performance. These involve, respectively, ensuring that the tested problem: (i) compiles without warning; (ii) gives the same results as the previous test; (iii) has not eroded significantly in performance. A specific test is defined by a problem configuration and a benchmark. As developers add new modules, they are required to add corresponding test implementations to the testing suite. FLASH_test then runs nightly on a large and growing list of platform/compiler combinations. A summary of results is published both internally for our developers and via a public web server to the community at large (<http://flash.uchicago.edu>). When problems arise, possible culprits are automatically notified, and a more or less well-defined process is carried out until the issues are resolved. Figure 1 shows a snapshot of the test suite results page. The window on the left side of the image is the main page, which shows the tests by date performed on a variety of computing platforms. A green or red square next to the date indicates pass or fail, respectively. Moving the cursor over the test date produces a pop-up window with the number of tests passed and the tests failed (if there are failures). Clicking the left mouse button with the cursor on the date produces a window with a listing of the tests run on that date, which in the figure is the window in the lower center of the image. From that window, one may click the left mouse button on a particular

test for more information, as is shown in the window on the right of the image.

The key point is that this process assumes verification as its starting point and automates and formalizes the re-verification process. Technically, this is a deceptively complicated process to carry out smoothly. Equally important, though, are several subtle strategic decisions that were glossed over above. For example, what is meant by “giving the same results as the previous test”? Should we require that the results match to machine precision point-by-point, within roundoff error, for certain global statistics? Are these both inter- and intra-machine comparisons? Furthermore, there is obviously no such thing as a complete set of tests. By what guidelines is the line drawn?

These are deep questions to which there seem to be no simple, general answers. We only comment on how they are addressed in our day-to-day software development practices. Practically speaking, we require every module in the code to appear in at least one test, and in theory for each dimensionality and a range of other obvious parameters (we still fall short in some cases). Our comparison testing methodology is to automate as much as possible while taking the most conservative approach to finding errors. Thus, our regression tests are done point by point and to machine precision. In some cases this causes false alarms (e.g. change in order of floating point operations), though such episodes are less common than we initially anticipated. The test is then flagged as a failure and the test “owner”, the researcher who designed the test, is notified to study the failures. The test owner can indicate via a web interface that research is being done on the possible failure to both alert other developers and instruct FLASH_test to halt subsequent warnings. The testing facility then provides the developer with some of the information necessary to verify the new solution (e.g. global quantities, comparison to known solution, etc.). However, we emphasize that the problem owner must redo his own original verification process. While the test suite provides as much information as possible to assist in this process, it is not automatic.

Though necessarily incomplete, the above process defines our criterion for re-verifying FLASH – that is, determining it is still as correct today it was correct yesterday. It is our experience that, without such a tool, group development of scientific code would be nearly impossible. At the very least, tremendous amount of time would be wasted, and many unnecessary bugs would be introduced beyond what might have been missed by the verification process itself.

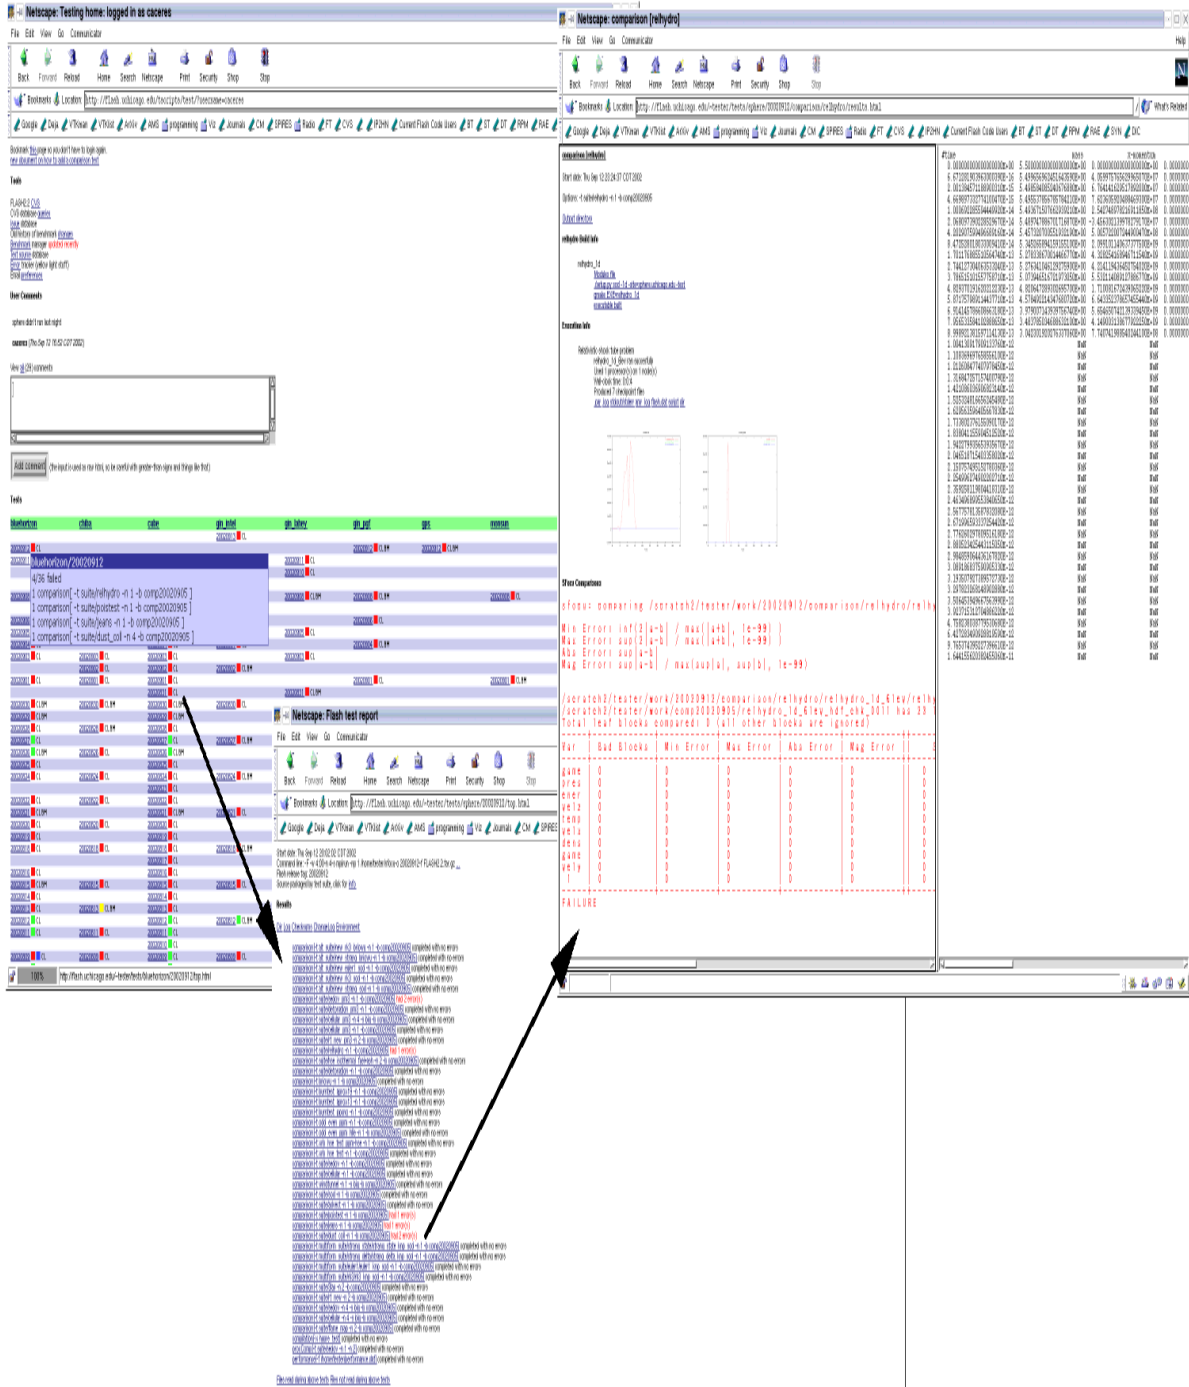


Fig. 1.— Snapshots from the test suite web page illustrating the “cascade” of information.

1.4.3. *Validation*

Validation is the process of determining the degree to which a model is an accurate representation of the real world from the perspective of the intended uses of the model (AIAA 1998). Validating a simulation or model requires identifying the key elements involved and for each element (as well as the integrated code) constructing test problems that have the results of laboratory experiments as the accepted results. Key elements include parts of the code that describe the fundamental physical processes and, for our interests, include items such as transitions to low and high Mach numbers, transport of energy by conduction or radiation, transport of energy by advection (convection), source terms (e.g. nuclear burning), equations of state, opacities, and microscopic transport such as molecular diffusion and viscosity. Validation problems tend to be much more complex than verification problems and typically have no analytic solutions. These problems nevertheless involve phenomena that are sufficiently simple to be studied both by experiment and simulation and to allow a detailed, quantitative comparison of the results. In addition, validation testing includes systematic grid sensitivity studies to assess grid convergence error and assess the level of refinement necessary to capture the key physical effects (AIAA 1998).

Validating an astrophysical simulation code is challenging because astrophysical events are often complex and involve many interacting physical processes, each of which must be tested. In addition, there are difficulties in making observations. The interiors of stars, for instance, do not lend themselves to direct observation, and, even if this were not a problem, the vast majority of astrophysical objects are too far away to resolve. Observations of thermonuclear flashes can only show the results of the events (light curves and spectra) and not the details of initiation of the outburst. Also, the length scales of the astrophysical objects present challenges. Flows within stars, for example, are expected to occur at Reynolds numbers greater than 10^9 , and terrestrial experiments cannot approach such a regime.

Another issue worth mentioning is the choice of equation of state. Phenomena such as degeneracy and other interactions play a critical role in astrophysics, and are expressed in the simulation through the equation of state. The properties of an equation of state (e.g. convexity) are particularly important when using hydrodynamics methods that solve the Riemann problem as does the principal hydrodynamics module in FLASH (Menikoff & Plohr 1989). Considerable effort can go into validating a hydrodynamics code with a known, accurate equation of state, but coupling it to an inappropriate or inaccurate equation of state will invalidate any simulations. Also, effort expended in testing equations of state relevant for a particular laboratory experiment will not necessarily improve confidence in predictions made by astrophysical simulations that require another equation of state.

Because of these complexities, it seems unlikely that one could satisfactorily validate an

astrophysical simulation. Instead, astrophysical validation efforts must focus on laboratory experiments that capture the relevant physics, with the expectation that experience gained from these closely related cases builds confidence in the predictions of the astrophysical simulations. Accordingly, a significant part of the challenge of validating astrophysical simulation codes is to find acceptable laboratory experiments. A good experiment for validation should be a good experiment itself (that is, it should provide accepted, repeatable results), it should adequately capture a significant portion of the physical processes of interest, and it should be diagnosed well enough for a meaningful comparison to simulation.

Our methodology for validation is to apply all that is described above as best we can. We are a collaboration between experimentalists and theorists working toward understanding the same types of problems, and we use results from the experimentalists in the collaboration for the study. For each experiment we determined the quantities that could most reliably be measured (in the experiment and simulation) or the quantities of greatest interest, and compared those measurements to the equivalent results from simulations. We performed resolution studies of the simulations, and investigated how the agreement between the simulations and experiments changed with resolution. Where possible, we investigated results for both experiments and simulations that were predicted by theory.

We note that our approach may be called “piecemeal” V&V. That is, we look carefully at each part of the simulations or models, testing and comparing to experiment. This approach to V&V should allow us to study regions of applicability for each component of the models or simulations, which in turn will provide insight into the applicability of integrated simulations of the entire physical phenomena. We are aware, however, that the reliability of integrated simulations based on separately validated components is suspect because interactions between the components may introduce additional sources of error. Where possible, integrated simulations should be validated in their own right, but this is difficult in astrophysics.

1.4.4. *Calibration*

A final subject to describe in the methodology of V&V is calibration. Calibration is not validation. Instead, calibration is a process performed in order to improve the agreement of computational results with experiments, and calibration does not generate the same level of predictive confidence as validation. Calibration is performed when there is uncertainty in the modeling of complex processes and also when there are incomplete or imprecise measurements in the experiments. Calibration involves adjustments to parameters in subgrid models, reaction rates, and boundary conditions, and it includes assumptions

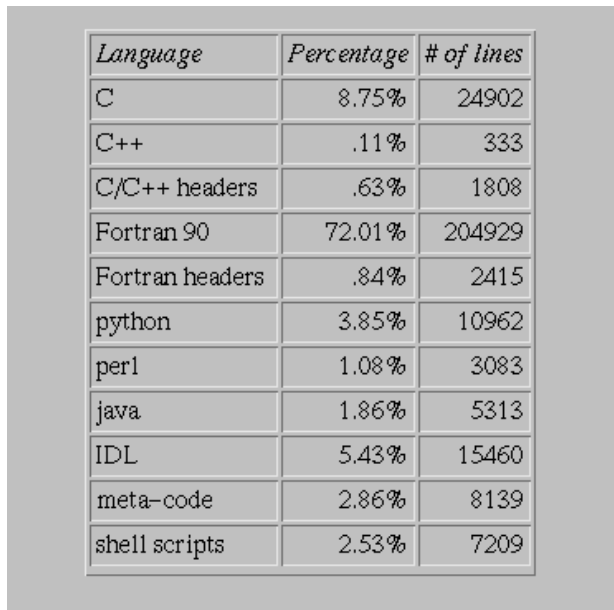
about minimal or optimal levels of mesh refinement that are made in cases where there are not completely resolved solutions (AIAA 1998). We note that because of the difficulty of validating simulations of astrophysical events, much of our effort may be considered calibration.

2. Numerical Method

The FLASH code is a parallel, adaptive-mesh simulation code for studying multi-dimensional compressible reactive flows in astrophysical environments. It runs on single- and multi- processor commodity PC's and workstations, on which the majority of development and debugging is executed; however, it is designed for efficiency on parallel supercomputers, and astrophysical simulations have been carried out on several of the top 10 machines on the Top500 (2002) list. In one test, FLASH achieved a sustained performance of 238 GFLOPS on 6420 processors on ASCI Red, for which it was awarded the SC2000 Gordon Bell prize, special category (Calder et al. 2000). Run times can last from seconds to days, depending mainly on the purpose of the simulation. Parallel processing is achieved using the Message-Passing Interface library (Gropp, Lusk, & Skjellum 1999).

One way efficiency is gained is by adaptive mesh refinement (AMR), in which resolution elements are added areas of complex flow and removed from quiescent regions. Spatial AMR capability is provided by a customized version of the PARAMESH library (MacNeice et al. 1999, 2000). The current models used for simulations assume that the flow is described by the Euler equations for compressible, inviscid flow. FLASH regularizes and solves these equations by an explicit, directionally split method (described below), carrying a separate advection equation for the partial density of each chemical or nuclear species as required for reactive flows. The code does not explicitly track interfaces between fluids so some numerical mixing can be expected during the course of a calculation.

At this writing, FLASH consists of 229,000 lines of source code and is implemented mostly in Fortran 90. This count includes the PARAMESH source code. Figure 2, a snapshot from the code statistics web page, shows the breakdown of the languages used in the code. FLASH emulates modern object-oriented software technology that allows for minimal effort to swap or add physics modules. Accordingly, the development of FLASH requires development and testing of each module as well as development and testing of the framework integrating the modules. In the subsection below, we provide details of the principal hydrodynamics module in FLASH. Complete details concerning the algorithms used in the code, the structure of the code, selected verification tests, and performance may be found in Fryxell et al. (2000) and Calder et al. (2000).



<i>Language</i>	<i>Percentage</i>	<i># of lines</i>
C	8.75%	24902
C++	.11%	333
C/C++ headers	.63%	1808
Fortran 90	72.01%	204929
Fortran headers	.84%	2415
python	3.85%	10962
perl	1.08%	3083
java	1.86%	5313
IDL	5.43%	15460
meta-code	2.86%	8139
shell scripts	2.53%	7209

Fig. 2.— Chart showing the proportions and numbers of lines of code for the languages used in FLASH.

2.1. Hydrodynamic Module

The primary hydrodynamic module in FLASH is based on the PROMETHEUS code (Fryxell, Müller, & Arnett 1989) and evolves systems described by the Euler equations for compressible gas dynamics in one, two, or three dimensions. The evolution equations are solved using a modified version of the Piecewise-Parabolic Method (PPM), which is described in detail in Woodward & Colella (1984) and Colella & Woodward (1984). PPM is a shock capturing scheme in which dissipation is used to regularize the Euler equations. (See Majda (1984) for a discussion of the importance of dissipative mechanisms.) Modifications to the method include the capability to use general equations of state (Colella & Glaz 1985). PPM is a higher-order version of the method developed by Godunov (1959, 1961). Godunov methods are finite-volume conservation schemes that solve the Riemann problem at the interfaces of the control volumes to compute fluxes into each volume. The conserved fluid quantities are treated as cell averages that are updated by the fluxes at the interfaces. This treatment has the effect of introducing explicit non-linearity into the difference equations and permits the calculation of sharp shock fronts and contact discontinuities without introducing significant non-physical oscillations into the flow. The original Godunov method is limited to first-order accuracy in both space and time because the distribution of each variable in each control volume is assumed to be constant. PPM extends this method by representing the flow variables as piecewise-parabolic functions and also by incorporating monotonicity

constraints to limit unphysical oscillations in the flow. PPM is formally accurate to only second order in both space and time, but performs the most critical steps to third- or fourth-order accuracy. This results in a method which is considerably more accurate and efficient than most second-order codes using typical grid sizes. A fully third-order (in space) method provides only a slight additional improvement in accuracy but results in a significant increase in the computational cost of the method.

PPM is particularly well-suited to flows involving discontinuities such as shocks and contact discontinuities. The method also performs well for smooth flows, although other schemes that do not perform the additional steps for the treatment of discontinuities are more efficient in these cases. The high resolution and accuracy of PPM are obtained by the explicit non-linearity of the scheme and through the use of smart dissipation algorithms, which are considerably more effective at stabilizing shock waves than the more traditional explicit artificial viscosity approach. Typically, shocks are spread over only one to two grid points, and post-shock oscillations are virtually nonexistent in most cases. Contact discontinuities and interfaces between different fluids create special problems for Eulerian hydrodynamics codes. Unlike shocks, which contain a self-steepening mechanism, contact discontinuities spread diffusively during a calculation; they continue to broaden as the calculation progresses. PPM contains an algorithm that prevents contact discontinuities from spreading more than one to two grid points, no matter how far they propagate.

The PPM implementation in FLASH is a directionally split, Direct Eulerian formulation. The hydrodynamics module applies the PPM method in one-dimensional sweeps across a block of data, advancing the time two steps. Reversing the order of the sweep for the second time step preserves second-order accuracy in time (Strang 1968). In three dimensions, the sweeps are performed in the order $xyz - zyx$ (for Cartesian geometry). The algorithm uses a nine-point stencil in each direction, requiring that each block have four ghost zones on each side. In addition, a small multi-dimensional artificial viscosity is added to provide a weak coupling between adjacent rows and columns in the directionally split scheme.

3. A Verification Test

We present a verification test problem that typifies the tests we have performed with the hydrodynamics module. The test problem is the advection of an isentropic vortex (Shu 1998; Yee, Vinokur, & Djomehri 2000). This two-dimensional problem exposes the directional splitting of the hydrodynamics algorithm to scrutiny. This problem was chosen because its solution is smooth and an exact solution is available. The simulations had constant time steps at each resolution such that (i) each simulation ended at exactly the same evolution

time, and (ii) the ratio of $\Delta t/\Delta x$ was constant across different resolutions, corresponding to a fixed Courant number.




Fig. 3.— Plot of L2 density error vs. mesh spacing for a two-dimensional advection test consisting of an isentropic vortex propagating diagonally across the simulation mesh. Shown are results from two sets of simulations, one with full PPM and one with non-linear steepening turned off. A power-law fit to the curves gave exponents of 2.13 for the default curve and 2.03 for the discontinuity-free curve.

In our test simulations, the vortex propagates diagonally with respect to the grid. The simulation domain is a square, $-5.0 \leq x, y \leq 5.0$. The ambient conditions are, in non-dimensional units, $u_\infty = 1.0$, $v_\infty = 1.0$, and $T_\infty = 1.0$. The following perturbations are added to get u , v , and T fields:

$$\delta u = -y \frac{\beta}{2\pi} \exp\left(\frac{1-r^2}{2}\right) \quad (1)$$

$$\delta v = x \frac{\beta}{2\pi} \exp\left(\frac{1-r^2}{2}\right) \quad (2)$$

$$\delta T = -\frac{(\gamma-1)\beta}{8\gamma\pi^2} \exp(1-r^2) \quad (3)$$

where $r^2 = x^2 + y^2$, $\gamma = 1.4$ is the ratio of specific heats and $\beta = 5.0$ is a measure of the

vortex strength. The density is then computed by $\rho = (T_\infty + \delta T)^{1/(\gamma-1)}$. The conserved variables (density, x - and y -momentum, and total energy) can be computed from the above quantities. The flow field is initialized by computing cell averages of the conserved variables: each average is approximated by averaging over 10^2 subintervals in the cell.

For all isentropic vortex simulations, periodic boundary conditions were applied, the time step was fixed and the Courant number was approximately 0.94, and the error was calculated at time $t = 2.0$. Solutions were obtained on 40^2 , 80^2 , 160^2 , 320^2 , 640^2 , and 1280^2 equispaced meshes. The exact solution at time t is the initial condition translated by $(u_\infty t, v_\infty t)$; it was computed by applying the same steps as for the initialization but with the vortex center translated by the appropriate amount, so that the cell-averaging of the conserved variables was consistent between the exact solution and the initial condition.

Figure 3 shows the L_2 error in density is plotted vs. the mesh spacing for two cases of the isentropic vortex problem. The “default” case was computed using the code as it is most often used for production simulations: no tuning of PPM parameters or adjustments to the algorithm were made. In the “discontinuity-free” case, several non-linear components of the PPM algorithm that improve behavior at shocks and contacts were disabled. Contact steepening, shock flattening, monotonization, and artificial viscosity are designed to improve or stabilize computations that contain discontinuous flow features, at the possible expense of increasing local truncation errors and reducing the convergence rate. Since the solution of the isentropic vortex problem is smooth these components are not required, and by comparing the discontinuity-free case to the default case, their influence can be examined.

The figure shows that the error decreases as expected as the mesh is refined; power law fits give convergence rates of 2.13 for the default simulation and 2.03 for the discontinuity-free simulation. The error is slightly higher for the default case on coarser grids, but on the finer grids the default and discontinuity-free versions of the code produce essentially identical results. We have varied many aspects of these simulation to see their effects on the error and the convergence rate. We found that while the measured errors varied, the convergence rates were not sensitive to the details of the initialization process (point-values, interpolation, number of subintervals for estimating cell-averages), to smaller time steps (and Courant numbers), or to whether or not the vortex was advecting. In all cases the code demonstrated second-order convergence.

4. Validation Tests

In the following sections, we present the results of validating FLASH with two laboratory experiments. These efforts focus only on validating the principal hydrodynamics module in FLASH, and the validation of other code modules such as burning and gravity is the subject of ongoing research. Where possible, we have quantified the results of the simulation-experiment comparison. Because the intent of this work was validation of the hydrodynamics module in FLASH and because the difficult problem of validating a material equation of state for terrestrial materials is beyond the scope of our efforts, the simulations presented below made use of simple gamma-law equations of state.

The validation tests presented below were performed on an adaptive mesh. In the case of the laser-driven shock simulations, the standard criteria for mesh refinement (testing the magnitude of the second derivative of density and pressure and refining or de-refining the mesh in regions where the magnitude is above or below a threshold) worked well to capture the shocks and discontinuities of the flow. In the Rayleigh-Taylor simulations, to avoid the possibility of under-resolving the initial conditions, the simulation domain in the region of the initial perturbations was forced to be fully refined. Beyond that region, the simulation applied the standard criteria for mesh refinement.

4.1. Laser-driven Shock Simulations

Intense lasers offer the chance to probe experimentally high energy density environments similar to those that exist in complex astrophysical phenomena. Such experiments are obvious choices for code validation. Holmes et al. (1999) performed a careful study of such experiments, investigating the Richtmyer-Meshkov (Richtmyer 1960; Meshkov 1969) instability. The study included experimental, numerical, and theoretical work and produced a quantitative comparison between results. Our efforts focus on modeling experiments performed using the Omega laser facility at the University of Rochester (Soures et al. 1996; Boehly 1995; Bradley 1998) that involve shock propagation through a multi-layer target. These experiments are designed to replicate the hydrodynamic instabilities thought to arise during supernova explosions. In addition to validation, the experiments may provide a better understanding of the turbulent mixing that occurs as a result of instabilities driven by the propagation of a shock through a layered target.

The experiment we used consists of a strong shock driven through a target with three layers of decreasing density. The interface between the first two layers is rippled while the second interface is flat. The planar shock is perturbed as it crosses the first interface and




Fig. 4.— Results of the 3-layer target experiment. Shown are side-on X-ray radiographs at 39.9 ns (left) and 66.0 ns (right). The long, dark “fingers” are spikes of expanding Cu, and the horizontal band of opaque material to the right of the spikes of Cu is the brominated plastic tracer showing the imprinted instability growth at the plastic-foam interface.

excites a Richtmyer-Meshkov instability. The perturbed shock then propagates through the second interface, imprinting the perturbation on the interface and leading to the growth of additional fluid instabilities. This three-layer experiment is meant to model the configuration of a core collapse supernova. In this case, it has been proposed that the development of fluid instabilities followed by mixing is responsible for certain features present in spectra obtained during the first few hundred days after the explosion (e.g. Arnett, Fryxell, & Müller 1989). The accepted scenario involves a supernova shock propagating through the outer layers of the star, which is composed of shells of different chemical compositions. The interaction of the shock with the interfaces between these shells leads to the development and growth of Richtmyer-Meshkov and Rayleigh-Taylor instabilities. The effect of these instabilities would be mixing of the material in the core of the star with material in the outer regions, which may be able to explain the early observation of radioactive core elements in SN 1987A (Kifonidis et al. 2000, and references therein).

The target consists of three main layers of material in a cylindrical Be shock tube, with the initial density decreasing in the direction of shock propagation. The materials are Cu, polyimide plastic, and carbonized resorcinol formaldehyde (CRF) foam, with thicknesses of 85, 150, and 1500 μm and densities 8.93, 1.41, and 0.1 g cm^{-3} , respectively. Performing the experiment with the target inside a shock tube delays the lateral decompression of the target, giving a more planar shock. Be is chosen as the material for the shock tube as it is essentially transparent to the diagnostic X-rays. The surface of the Cu layer is machined with a sinusoidal ripple of wavelength 200 μm and amplitude 15 μm . The laser drive end of the target consists of a 10 μm section of CH ablator to prevent direct illumination of the target and the associated pre-heating of the rest of the target. Embedded within the

polyimide layer is a $75\text{ }\mu\text{m}$ thick, $200\text{ }\mu\text{m}$ wide (along the diagnostic line of sight) tracer strip of brominated CH (4.3% by number of atoms in the material, i.e. the atomic composition is $\text{C}_{500}\text{H}_{457}\text{Br}_{43}$).




Fig. 5.— Schematic of the 3-layer target simulation initial conditions. Shown are the locations of the three materials, Cu, CH, and C, the shock, and the details of the sinusoidal perturbation of the Cu-CH interface. The schematic is not to scale.

The experiment is driven by 10 beams of the Omega laser with a nominal measured energy of 420 J/beam in a 1 ns pulse at a laser wavelength of $\lambda_L = 0.351\text{ }\mu\text{m}$. The peak intensity (in the overlapped spot) is $7.2 \times 10^{14}\text{ W cm}^{-2}$, while the average intensity is $5.7 \times 10^{14}\text{ W cm}^{-2}$. The shock is perturbed as it crosses the corrugated Cu-polyimide interface and oscillates as it propagates through the polyimide/CH(Br). When it reaches the foam interface, it imprints the perturbation. The experiment is observed side-on with hard X-ray radiography using a gated framing camera. Eight additional laser beams are focused on an iron back-lighter foil located near the target and generate 6.7 keV X-rays to which the Cu and CH(Br) tracer strip are opaque and the polyimide and foam are nearly transparent. Nearly all of the contrast at the polyimide/CH(Br)-foam interface comes from the tracer layer. This allows visualization of the shock-imprinted structure at that interface over only the central $200\text{ }\mu\text{m}$ of the target along the line of sight without edge effects near the wall of the shock tube. Full details of the experiment may be found in Kane et al. (2001) and Robey et al. (2001).

A recent study by Robey et al. (2002) addressing the issue of the onset of turbulence in laser-driven shock experiments provides an estimate of the Reynolds numbers for these experiments. The study notes that recent experimental work by Dimotakis (2000) indicates that for a wide range of stationary flow geometries, there appears to be a nearly universal value of Reynolds number ($1\text{--}2 \times 10^4$) at which an abrupt transition to a well-mixed state

occurs, and this transition has been suggested as an indicator for the transition to fully-developed turbulence. By combining experimental measurements with a kinematic viscosity estimate, the Robey et al. study indicates that the time dependent Reynolds number can approach 10^5 during the laser-driven shock experiments and that though there are some caveats, the experiments are “perhaps very close but somewhat short of the threshold value required for the onset of a mixing transition.” This would indicate that these sorts of experiments may be approaching the transition to fully-developed turbulence. We note that the laser-driven experiment described in this study differs somewhat from the laser experiment of our validation test, but it demonstrates the flow regimes that laser-driven shock experiments can reach.




Fig. 6.— Gray scale images of CH abundance at approximately the time of the late time experimental image, 66.0 ns, from simulations at varying resolutions. The effective simulation resolutions were, top to bottom on left followed by top to bottom on right, 128×64 , 256×128 , 512×256 , 1024×512 , 2048×1024 , 4096×2048 , corresponding to 4, 5, 6, 7, 8, and 9 levels of adaptive mesh refinement.

Figure 4 shows the results of the 3-layer target experiment. The images are X-ray radiographs at two times, 39.9 ns (left) and 66.0 ns (right). The long, dark “fingers” are spikes of expanding Cu, and the horizontal band of opaque material to the right of the spikes of Cu is the brominated plastic tracer, showing the imprinted instability growth at the plastic-foam interface. The length of the Cu spikes in the experiment was determined by

three methods. The first method was a straightforward visual inspection of the images using as a spatial reference a gold grid of $63.5\ \mu\text{m}$ period, located just below the images of Figure 4. The second method used a contour routine to try to better quantify the uncertainty in the location of the edges of the spikes. The third method was done in a manner consistent with the analysis of the numerical simulations. A $500\ \mu\text{m}$ section in the center of the images was vertically averaged to produce a single spatial lineout of optical depth through the region occupied by the Cu and CH. The same 5% and 90% threshold values as the simulation analysis were used to quantitatively determine the extent of the Cu spikes. Taking the average of all three methods, values of $330 \pm 25\ \mu\text{m}$ and $554 \pm 25\ \mu\text{m}$ are obtained at 39.9 and 66.0 ns, respectively.

There are several sources contributing to error in these experimental measurements: the spatial resolution of the diagnostic, the photon statistics of the image, target alignment and parallax, and the specific contrast level chosen as the definition of the length of the Cu spikes. The intrinsic diagnostic resolution is set by the imaging pinhole, which is $20\ \mu\text{m}$ in diameter. The photon noise statistics of the image can further degrade the spatial resolution when the number of photons per detector resolution element is small (see recent reference: Landen, et al. 2001). For the present number of backlighter beams, diagnostic magnification, and pinhole size, the photon statistics produce a signal-to-noise ratio > 20 , and therefore do not contribute to a further decrease in spatial resolution. The target alignment with respect to the diagnostic line-of-sight is generally within 1%. Since most of the contrast comes from the $200\ \mu\text{m}$ wide radiographic CH(Br) tracer layer, this contributes an additional 3.5% to the spatial uncertainty. Finally, the specific contrast level chosen to represent the spatial extent of the spikes as discussed above contributes to the uncertainty. This uncertainty was estimated from the variation resulting from the three methods used to determine the spike lengths.

In addition to the spatial error, there are also several sources of uncertainty in the temporal accuracy of the measurements. These arise from target-to-target dimensional variations, shot-to-shot drive intensity variations, and the intrinsic timing accuracy of the diagnostics. The nominal target dimensions used in the simulations were $10\ \mu\text{m}$ of CH (ablator layer), $85\ \mu\text{m}$ Cu, $150\ \mu\text{m}$ of polyimide CH(4.3% Br), and $1500\ \mu\text{m}$ of CRF. In the actual targets, however, there is some variation from these nominal values, and this variation enters into the quantification of temporal error between experiment and simulations. For the target that produced the image at 39.9 ns, the dimensions of the 4 layers are 10, 105, 145, and $1505\ \mu\text{m}$. For the target that produced the image at 66.0 ns, the dimensions of the 4 layers are 10, 98, 150, and $1514\ \mu\text{m}$. The biggest difference between the experimental target dimensions and those used in the simulations is in the dense Cu layer, where the experimental targets are $\sim 20\%$ thicker. Shock propagation through these slightly different




Fig. 7.— Percent difference of the Cu spike lengths from those of the highest resolution (9 levels of adaptive mesh refinement) simulation vs. time. The percent differences are from the lower resolution simulations of 4, 5, 6, 7, and 8 levels of adaptive mesh refinement. We note that the convergence is not perfect. The curve from the 8 level of refinement simulation crosses those of the 6 and 7 level of refinement simulations, indicating a higher percent difference.

thicknesses will cause a small discrepancy in the timing.

The variation in the drive intensity from shot-to-shot is another possible source of temporal uncertainty. For the image obtained at 39.9 ns, the average laser drive energy was 416 J(UV)/beam. For the shot at 66.0 ns, the drive energy was 421.5 J(UV)/beam. Therefore, there is a 1.3% intensity variation. The drive pressure scales as the intensity to the 2/3 power, and material velocities scale as the square root of the pressure (Lindl 1998), so the velocity variations from shot-to-shot will be less than 0.5%. Shot-to-shot timing variations resulting from this mechanism will be of the same order, i.e. less than one percent. The intrinsic temporal accuracy of the diagnostics is even better than this, with a typical uncertainty of ± 200 ps. The experimental uncertainty in the timing is therefore relatively small, and is approximately indicated by the width of the symbols used in the figure (described below) that compares the experimental results to the simulation results.

The simulation models the experiment with a similar three-layer arrangement. The




Fig. 8.— Cu spike length vs. adaptive mesh refinement level. Shown are the spike lengths at the two times corresponding to those of the experimental results from all of the simulations. The percent differences shown in Figure 7 were calculated at many times in the simulations. This figure illustrates how the spike length changed with resolution at the two times of the experimental results.

three materials were Cu, polyimide CH, and C with the same densities as the three layers of the actual target. The simulation began 2.1 ns into the experiment, at which point the shock is approaching the Cu-CH interface. The initial thermodynamic profiles were obtained from simulations of the laser-material interaction performed with a one-dimensional radiation hydrodynamics code (Larson & Lane 1994). The results were mapped onto the two-dimensional grid with a perturbed Cu-CH interface at a simulation time of 2.1 ns and were then evolved out to approximately 66 ns. The materials were modeled as gamma-law gases, with $\gamma = 2.0$, 2.0, and 1.3 for the Cu, CH, C, respectively. These values for gamma were chosen to give similar shock speeds to those observed in the experiments. Figure 5 illustrates the initial configuration. The simulation used periodic boundary conditions on the transverse boundaries and zero-gradient outflow boundary conditions on the boundaries in the direction of the shock propagation. Note that use of periodic boundary conditions in the transverse directions is not in keeping with the boundary conditions of the experiment. The experiment was performed with the three materials of the target inside a cylindrical Be




Fig. 9.— Simulated radiographs from the six level of refinement (effective resolution of 512×256) simulation of the three-layer target experiment. The simulated radiographs were created from the fluid abundances at times corresponding approximately to those of the images from the experiment, 39.9 ns (left) and 66.0 ns (right). Shown are the parts of the simulation domain that match the regions in the experimental results.

shock tube. Accordingly, the experiment results show the influence of the shock tube walls as a curving or pinching of the outer Cu spikes. Our simplified model did not consider these boundary effects.

We first present the results of a resolution study of the simulation. We performed six simulations from equivalent initial conditions, changing only the maximum mesh resolution. The results of the six simulations are illustrated in Figure 6. As described above, FLASH solves an advection equation for each abundance, allowing us to track the flow of each material with time. Shown are fluid abundances for the CH, the intermediate material, at approximately 66 ns. The abundance is represented by a gray scale, so that the white regions (for which the abundance is zero) to the left and right of the central gray region are Cu and C, respectively. Thus, the CH abundance shows features of both material interfaces, and its evolution shows the behavior of both the spikes of Cu and bubbles of C. The effective simulation resolutions were, top to bottom on left then top to top to bottom on right, 128×64 , 256×128 , 512×256 , 1024×512 , 2048×1024 , 4096×2048 , corresponding to 4, 5, 6, 7, 8, and 9 levels of adaptive mesh refinement.

We found that all of the simulations capture the expected bulk properties of the flow; that is, all showed similar spikes of Cu and bubbles of C. The principal differences in the results were in the amount of small scale structure in the flow, with the amount of small-scale structure increasing with resolution. The increasing amount of small scale structure in the simulations should be readily apparent in the panels of the figure. This behavior is expected because the dissipation mechanism in PPM operates on smaller and smaller scales as the

resolution is increased. We note that the increasing amount of small-scale structure seen with increasing resolution indicates that these simulations will not demonstrate a converged flow with higher resolution until the turbulent scales are fully resolved; only at this point will we be resolving all the relevant length scales in the problem.

Because, as noted above, we observe increasing amounts of small-scale structure in the simulations with higher resolutions, convergence studies should focus on integrated quantities such as instability growth, which may converge. In order to quantify the results, we examined the lengths of the Cu spikes, an integral quantity of the flow that we were best able to accurately measure both in the simulation results and in the experimental results. The spike lengths were measured by averaging the CH abundance in the y -direction across the simulation domain then smoothing the resulting one-dimensional array slightly to minimize differences that would occur owing to very small scale structure. The length of the Cu spikes was then determined by the average distance spanned by minimum locations of average abundances 0.05 and 0.9. The results were reasonably robust to the amount of smoothing and threshold values.

The results of testing the convergence of the Cu spike length measurements are shown in Figures 7 and 8. Figure 7 shows percent differences from the highest resolution simulation as functions of time. The trend is that the difference decreases with increasing mesh resolution, with the seven and eight level of adaptive mesh refinement simulations always demonstrating agreement to within five percent. The trend of decreasing difference with increasing mesh resolution demonstrates a convergence of the flow, but it is subject to caveats. We note that the trend does not describe the behavior at all points in time (that is, the percent difference curves sometimes cross each other), and this average measurement is an integral property of the flow and in no way quantifies the differences in small scale structure observed in the abundances. In particular, we note that the difference curve for the simulation with eight adaptive mesh refinement levels crosses the curves of both the seven and six level simulations, suggesting that higher resolution simulations may deviate further from these results and produce degraded agreement with the experiment. Figure 8 shows the Cu spike lengths at approximately the two times of the experimental results, 39.9 and 66.0 ns, vs. number of refinement levels, quantifying the changes of the Cu spike lengths at the times of the experimental results with resolution.

Figures 9 and 10 show simulated radiographs from two simulations. Both figures show a simulated radiograph at approximately the two times corresponding to the images from the experiment, 39.9 ns (left panel) and 66.0 ns (right panel). The simulation in Figure 9 had six levels of adaptive mesh refinement (an effective resolution of 512×256), while the simulation in Figure 10 had seven levels of refinement (1024×512). The radiographs were

created from the abundances of the three materials. An artificial opacity was assigned to each abundance, Cu having the highest, CH the intermediate, and C the lowest. Then the opacity was applied to intensity from an artificial “beam” that consisted of a circular region that linearly decayed above a certain radius from a point near the center of the simulation domain. In addition, the abundances were de-resolved to match the resolution of the pixels in the images from the experiment and random Poisson ‘noise’ was added to the intensity. Comparison of the simulated radiographs to the radiographs from the experiment show that the simulations captured the bulk behavior of the materials, particularly the growth of Cu spikes and the development of C bubbles.




Fig. 10.— Simulated radiographs from the seven level of refinement (effective resolution of 1024×512) simulation of the three-layer target experiment. The simulated radiographs were created from the fluid abundances at times corresponding approximately to those of the images from the experiment, 39.9 ns (left) and 66.0 ns (right). Shown are the parts of the simulation domain that match the regions in the experimental results.

It is worth noting that the increased amount of small scale structure in the higher resolution simulation is visible in comparing the radiographs even though the images have been de-resolved. From inspection, the amount of small scale structure in the six level simulation appears to agree with that observed in the experiment better than the seven level simulation. This agreement indicates that the amount of viscosity in the experiment is higher than the numerical viscosity of the higher resolution runs. Caution concerning this interpretation is warranted, however, because the experimental radiograph, which is essentially the shadow of the material, will almost certainly fail to capture the true amount of small scale structure. Conclusions regarding the correct amount of small scale structure (and the correct Reynolds number of the flow) must await better diagnostic resolution of the experiments followed by a careful, quantitative comparison with simulation results.

Figure 11 shows the Cu spike length vs. time for the six, seven, eight, and nine levels of adaptive mesh refinement simulations. Also shown in the figure are the experimental results.

The error bars correspond to $\pm 25 \mu\text{m}$, the spatial error of the experiment. The width of the symbols marking the experimental results indicates approximately the timing error. The figure shows that the simulations agree well with the experimental results. In addition, inspection of the evolution of the shock as it passes through the CH shows that the shock oscillates as expected from the linear theory of compressible fluids (Dyakov 1954; Freeman 1955; Landau & Lifshitz 1987) and is almost planar when it reaches the CH/C interface. Figure 12 shows the logarithm of density on the entire simulation grid at approximately the time of the late experimental result for the simulation with seven levels of adaptive mesh refinement. In this case, the image is fully resolved (to the resolution of the simulation).




Fig. 11.— Cu spike length vs. time. The curves are from simulations at 6, 7, 8, and 9 levels of refinement (effective resolutions of 512×256 , 1024×512 , 2048×1024 , 4096×2048), and the points with error bars are results from the experiment. The error bars represent $\pm 25 \mu\text{m}$, and the width of the symbols represents the timing error.

4.2. Rayleigh-Taylor Simulations

The classic Rayleigh-Taylor problem consists of a dense fluid on top of a light fluid in the presence of a gravitational acceleration. The initial configuration is in an unstable equilibrium, and any perturbation of the fluid interface leads to instability growth. In




Fig. 12.— Full resolution image of the log of density from the seven levels of adaptive mesh refinement simulation at approximately the time of the late time experimental result. The spikes of Cu are visible as the reddish-yellow ($\rho \sim 2 \text{ g cm}^{-3}$) fingers moving into the less dense ($\rho \sim 0.5 \text{ g cm}^{-3}$) CH. The bubbles of C are the dark green ($\rho \sim 0.2 \text{ g cm}^{-3}$) regions to the right of, and opposing, the spikes of Cu. The transition from compressed C (rightmost yellowish region) to uncompressed C (blue region on far right) marks the position of the shock, which shows a slight perturbation.

the linear regime, the instability growth rate is proportional to the square root of the perturbation wave number k (Chandrasekhar 1981). If the interface is sharp, the problem is mathematically ill-posed because the growth rate diverges with large k . The simulations we performed began with completely sharp interfaces, but numerical diffusion in the simulation regularizes the problem by increasing the width of the boundary and thereby bounding the growth rate. In experiments, the growth rate is bounded either because surface tension limits the maximum k or because physical diffusion widens the boundary (Duff, Harlow, & Hirt 1962; Faber 1995).

The validation test consisted of multi-mode simulations performed from a standard set of initial conditions, allowing for comparison with results of other research groups and experiments. Single-mode simulations allow for testing of convergence of the solutions and are meant to determine the minimum resolution per perturbation wavelength for a converged simulation. This case is of particular interest because it illustrates the difficulty of code validation even for apparently simple laboratory experiments.

In the case of a Rayleigh-Taylor configuration with a multi-mode perturbation, bubble and spike mergers and bubble/spike competition are thought to lead to an instability growth

according to a t^2 scaling law, which may be written as (Youngs 1994)

$$h_{b,s} = \alpha_{b,s} g A t^2, \quad (4)$$

where $h_{b,s}$ is the height of a bubble or spike, g is the acceleration due to gravity, and $A = (\rho_2 - \rho_1)/(\rho_2 + \rho_1)$ is the Atwood number. $\rho_{1,2}$ is the density of the lighter (heavier) fluid, and t is the time. The coefficient α is a measure of the rate of potential energy release. The initial conditions for the multi-mode simulations consisted of velocity perturbations corresponding to modes 32–64. These initial conditions were adapted from a standard set used by a consortium of researchers. This consortium, known as the Alpha Group, was formed in 1998 by Guy Dimonte for the purpose of investigating the validity of the t^2 scaling law, and, if it holds, to determine the value of α . The study undertaken by the consortium includes both code-to-code comparisons and comparison with experiment (Dimonte et al. 2002).

The experimental multi-mode Rayleigh-Taylor results come from studies performed by Schneider, Dimonte, & Remington (1998) and Dimonte & Schneider (2000). The experiments investigated the Rayleigh-Taylor instability over a range of density ratios using a variety of sustained and impulsive acceleration histories. The experiments consisted of a sealed plastic fluid container accelerated by the Linear Electric Motor (Dimonte et al. 1996). The diagnostics used for comparison were provided by laser-induced fluorescence, in which a dye is added to the system and excited by a laser beam focused into a sheet propagating upwards through the cell. Images captured with a gated charged-coupled device and 35 mm film cameras allow diagnosis of the internal structure of the mixing zone. Figure 13 shows laser-induced fluorescence images from an experiment with $A = 0.32$ and a nominally constant acceleration of ~ 68 times the acceleration of gravity on Earth. The images were captured at 25 and 44 ms. The bi-level images shown are made from intensity images by setting the values to zero below and unity above an intensity threshold. The threshold was chosen as $\sim 50\%$ to conserve the two fluid volumes. In the images, the lighter material is on the top and appears black, and the heavier material is on the bottom and appears white. Figure 14 shows the bubble and spike magnitudes from the experiment. The resulting values for α were 0.052 for the bubbles and 0.058 for the spikes. Dimonte and Schneider report $\alpha \sim 0.05$ from several experiments with a constant acceleration. These results are in good agreement with earlier experimental work by Youngs (1989).

There are many unknowns that may influence the results of these experiments. There may be unaccounted-for noise in the experiment that can change the actual initial conditions. Since the value of α is thought to be dependent on the power spectrum of the initial perturbation, long wavelength noise in the experiment could make α larger because these modes would dominate the dynamics in the nonlinear regime. This situation will occur if the




Fig. 13.— Experimental results from a multi-mode Rayleigh-Taylor experiment performed on the Linear Electric Motor. Shown are bi-level laser-induced fluorescence images from an experiment with Atwood number $A = 0.32$ at $t = 25$ ms (left) and 44 ms (right). The dense material ($\rho_2 = 1.43 \text{ g cm}^{-3}$) is on the bottom and appears white. The light material ($\rho_1 = 0.73 \text{ g cm}^{-3}$) is on the top and appears black. The direction of the acceleration of the experimental capsule was down, providing an effective upward acceleration. The width of the material shown in each panel was 6.2 cm .

long wavelength modes never reach the self-similar regime given the length and time scales of the experiment.

In addition, the diagnostics of the experiments may lead to spurious results. For example, because the laser-induced fluorescence method illuminates the mixing zone with a planar sheet of light, this diagnostic can lead to aliasing of long wavelength structures into short wavelength features in the images, thus affecting the interpretation of observed small-scale structures in the mixing zone. Also, because of the dynamic limits on diagnostic resolution, the formation of small-scale structure cannot be completely determined.

Results of our multi-mode simulations are shown in Figures 15 - 18. The simulations began from a Rayleigh-Taylor configuration with $A = \frac{1}{2}$ ($\rho_1 = 1 \text{ g cm}^{-3}$ and $\rho_2 = 3 \text{ g cm}^{-3}$) and with $g = 2 \text{ cm s}^{-2}$. In the simulations, gravity acts in the y -direction, and the simulation domain consisted of a $10 \times 20 \times 10 \text{ cm}$ region, with the fluid interface 10.625 cm above the bottom of the domain. The initial interface displacement perturbation was converted to a velocity perturbation using the linear theory of the Rayleigh-Taylor instability (Chandrasekhar 1981). In order to completely resolve these initial perturbations, the mesh was forced to be fully resolved in the perturbed region of the simulation domain. This region

was 4 cm tall and centered on the initial fluid interface.




Fig. 14.— Plot of distance vs. gAt^2 from a multi-mode Rayleigh-Taylor experiment performed on the Linear Electric Motor. Shown are the magnitudes of bubble height and spike depth as functions of the product of acceleration (g), Atwood number (A), and the time squared (t^2). The slope of each curve equals α , the rate coefficient. For this experiment, fitting straight lines to the curves produced $\alpha = 0.052$ and 0.058 for the bubbles and spikes, respectively.

Figure 15 shows bubble heights and spike depths for two three-dimensional multi-mode simulations. The effective resolutions of the two simulations in the (x, y, z) directions were $128 \times 256 \times 128$ and $256 \times 512 \times 256$. The top two curves are bubble heights, and the lower two curves are spike depths. The bubble heights and spike depths were measured by averaging the density in each x - z plane, and comparing the average density profile to the initial density profile, marking the points at which the average density deviated from the initial profile by more than 1%. Results were robust to reasonable changes in this threshold value. The distances shown were measured from the initial fluid interface, and the distance between the two curves from each simulation is the width of the mixing zone. The results show some differences between the two simulations, but both show a predominantly linear growth above a certain point.




Fig. 15.— Plot of distance vs. time from two three-dimensional multi-mode simulations. Shown are bubble heights (the two top curves) and spike depths (the two lower curves) as measured from the initial fluid interface. The simulations had effective resolutions of $128 \times 256 \times 128$ and $256 \times 512 \times 256$.

Figure 16 shows results from the higher resolution simulation. Shown are bubble and spike magnitudes plotted vs. gAt^2 from the higher resolution three-dimensional simulation. The slope of a linear fit to each curve gives α , the rate coefficient. The results for α , obtained from a linear fit to each curve, were 0.024 and 0.030 for the bubbles and spikes, respectively. If we neglect the first five seconds as being part of a different stage of evolution than the merger-dominated regime that gives the t^2 behavior and look at the slopes, we find lower results for α , 0.021 and 0.026.

Figure 17 allows for visual comparison of the structure of the mixing zone of the higher resolution simulation to that of the experiment. Because the multi-mode simulations were not designed to match the experiments, (i.e. different Atwood numbers, accelerations, and geometry), we may make only a qualitative comparison. Figure 17 shows bi-level images of the simulation at $t = 8.75$ s. (left panel) and $t = 15.5$ s. (right panel). The images were created from the heavy fluid abundance in the x - y plane at $z = 2.5$ cm by setting the values to zero below and unity above an abundance of 0.5. As with the experimental results, the dense fluid is on the bottom and appears white, and the light fluid is on the top




Fig. 16.— Plot of distance vs. gAt^2 from the higher resolution ($256 \times 512 \times 256$) three-dimensional simulation. Shown are the magnitudes of bubble height and spike depth as functions of the product of acceleration (g), Atwood number (A), and the time squared (t^2). The slope of each curve equals α , the rate coefficient. For this simulation, fitting straight lines to the entire curves produced $\alpha = 0.024$ and 0.030 for the bubbles and spikes, respectively. If the first five seconds of evolution are neglected to eliminate the parts of the curves with a rapidly changing slope, a straight line fit yields $\alpha = 0.021$ and 0.026 . The thin vertical line at $gAt^2 = 25$ marks $t = 5$ s.

and appears black. The early time was chosen to match the proportion of evolution as that of the early time in the experimental images. Figure 18 is a rendering of density from the higher resolution simulation. Shown is the mixing zone, with well-developed bubbles and spikes. The colors indicate lower density (red), intermediate density (yellow), and higher density (green). The higher and lower density material above and below the mixing zone is transparent. As illustrated in these figures, the higher resolution multi-mode simulation shows a very similar structure to the experiments. Our results indicate $\alpha \leq 0.03$, however, which is not in good agreement with the experiments and indicates the presence of some systematic error. Other research groups in the consortium using Eulerian hydrodynamics methods report similar results, although groups using other methods report different results. Results of consortium studies will appear in publications of the Alpha Group (Dimonte et al.

2002).




Fig. 17.— Bi-level cross sections of the multi-mode simulation at $t = 8.75$ s (left) and 15.5 s (right). The images were produced from cross sections of the heavy fluid abundance in the x - y plane at $z = 2.5$ cm. The images are shown from the same perspective as those from the experiment with the heavy fluid appearing white and the light fluid appearing black and an upward direction of the acceleration.

5. Results

In this manuscript, we presented some of the results of our efforts at verification and validation of FLASH, our astrophysical simulation code. We presented the results of a verification test and the results of two validation tests in which we were able to carefully compare our simulation results to experimental results.

The verification tests we performed for this study indicate that the PPM hydrodynamics module is performing as expected. As illustrated in the isentropic vortex test problem, the error showed the expected order of convergence and demonstrated the effect of the special algorithms in PPM for discontinuities. We note that this verification test addresses only the hydrodynamics module. Great care must be taken in the operator splitting when




Fig. 18.— Rendering of the mixing zone of the higher resolution multi-mode Rayleigh-Taylor simulation. Shown is density at a simulation time of 15.5 s. The colors indicate lower density (red), intermediate density (yellow), and higher density (green). Densities higher or lower than those occurring in the mixing zone are transparent. The initial perturbation consisted of modes 32-64, with an effective resolution of $256 \times 512 \times 256$.

coupling the hydrodynamics module to other modules such as gravity so as to not degrade the time accuracy. Coupling the hydrodynamics module to non-time-centered body forces, for example, will produce first-order convergence in time despite the expected second-order convergence of PPM. Additional verification tests are being performed with other modules coupled to the hydrodynamics and with a modified version of the PPM hydrodynamics module for maintaining hydrostatic equilibrium (Zingale et al. 2002).

The results of laser-driven shock validation test shows that we can capture the bulk properties of the flow, including the morphological properties of the resolvable structures. We observed the expected instabilities at the two material interfaces that grew with time during the course of the simulations. We performed a mesh resolution study that in general showed convergence of instability growth at the material interfaces. The resolution study did not completely demonstrate convergence, though, because the amount of small scale structure present in the simulations increased with resolution as may be seen by a visual inspection of the results (Figure 6). We attribute this increasing amount of small scale

structure to the fact that the amount of numerical dissipation in our hydrodynamics method decreases with increasing resolution.

Visual inspection of the experimental results and our simulated radiographs suggested that a simulation with six levels of adaptive mesh refinement (an effective resolution of 512×256) better matched the observed small scale structure. We note, though, that the correct amount of small scale structure in the experiments is probably not captured in the experimental diagnostics because of resolution limits and because the radiograph produces a two-dimensional shadowgraph of a three-dimensional experiment, averaging out small structures. Therefore, we draw no conclusions about the resolution required to capture the correct amount of small-scale structure.

Measurement of the lengths of the spikes allowed us to quantify the results of the simulations, and we found that simulations with seven and eight levels of adaptive mesh refinement agreed to within 5% of the highest resolution (nine levels) simulation at all times during the evolution (Figure 11). Complete convergence is prohibited by the growth in small-scale structure with increasing resolution. The spike lengths at intermediate resolutions matched the experimental results very well, falling within the experimental errors. We interpret this as an important calibration result that indicates that six and seven levels of adaptive mesh refinement are appropriate for two-dimensional simulations of these particular experiments. In addition, the simulations showed that the shock propagating through the corrugated material interface develops a perturbation that oscillates and becomes planar at the second interface, as expected from theory. We also observed the expected imprinting of the perturbation on this second interface.

For this validation problem, comparison between simulation and experiment is limited by the diagnostic resolution of the experiments and the sophistication of the physics models included in the simulations. Improvements to experimental diagnostics will be made in the next generation of laser experiments, and improvements to the simulations could be made by inclusion of a more realistic material equation of state, adding the walls of the shock tube, and modeling the laser-driven energy deposition process. Absent these improvements, we cannot conclude that the good agreement of our results with the experiment completely validates the simulations.

Our multi-mode Rayleigh-Taylor instability simulations, like the simulations of the laser experiment, show that the code results agree with the observed bulk properties of the flow: we observe a mixing zone that is very similar in structure to that of the experiment. Our simulations do not, however, agree well with the experimental rate coefficient, α ; our results are systematically lower. This difference could be due to limitations of either the simulations or the experiments. The experimental initial conditions, which strongly effect the growth

rate, are poorly characterized. Further, limitations on experimental diagnostics may distort measured growth rates and mixing.

The simulations, on the other hand, may have been under-resolved. In our low-resolution simulation, most of the power in the initial conditions was in modes which were resolved with only 2–4 grid points, which is likely far too low. Increasing the resolution to resolve these modes with 4–8 points increases the observed growth rate, in (somewhat) closer agreement with experimental results.

5.1. Conclusions and Lessons Learned

The principal conclusions we draw from our efforts are that validation of an astrophysical simulation code is a difficult process and that the process led us to unanticipated questions. The verification tests we performed allowed for a careful quantitative study of the accuracy and convergence rates of FLASH. The exercise proved useful in a variety of ways, not the least of which was developing techniques for accurately comparing the results of our simulations to accepted answers. The validation tests we performed, while they did not conclusively validate the simulations, provided insight into the many issues involved in numerical modeling, and served to calibrate the code and build confidence in our results.

Validation testing allows for an assessment of error in the code and progress in building confidence in the results. We learned that the process is limited, however. Analytic solutions are typically available only for very simple problems that may not serve as strenuous tests of a particular code. Convergence studies, though essential, do not imply that the converged answer is correct. Code-to-code comparisons are another useful tool; they can be more probing than analytic solutions, and can shed insight into the behavior of different methods. But again, similar answers do not imply correctness.

Comparing numerical and experimental results is also difficult. Nonlinear systems typically have exponential sensitivity to initial conditions, so that any unmodeled initial perturbation in the system can greatly affect the results. We suspect this sensitivity to initial conditions may have played a role in the disagreement between the simulations and the experiment for the multi-mode Rayleigh-Taylor instability. Experimental results may be limited, noisy measurements of poorly-characterized but real physical situations, while the numerical results may be complete diagnostics of a situation derived from a model of dubious applicability. Comparisons must therefore rely on choosing features which are obtainable from both results, yet sensitive enough to the physics of interest that comparing the features is a strong probe of the model and numerics. The complexity of this procedure means that

doing it well requires feedback between the computational and experimental researchers, with all parties improving techniques and performing multiple runs. The computational simulations may be very expensive, however, and performing relevant experiments can be a costly and difficult process.

As our stated goal in this work was validating an astrophysical simulation code, the assumptions in the models and their applicability to both terrestrial and astrophysical flows warrants discussion. Our numerical models did not have an explicit model for viscosity. We saw, however, that the effects of numerical dissipation cannot be avoided in simulations. Other physical processes that may play a role but are not included in the models include surface tension, species diffusion, and thermal diffusion. With only Euler’s equations, it is impossible to adjust the relevant dimensionless numbers (Schmidt, Reynolds, Rayleigh, and Prandtl numbers, for example) to match the correct properties of the fluid.

Our use of models without an explicit viscosity resulted in an increasing amount of small scale structure observed with increasing resolution in simulations of both validation problems. As mentioned above, limits on diagnostic resolution prevent us from determining the correct amount of small scale structure in both validation experiments. Were the experimental diagnostics better, we could incorporate physically-motivated terms for viscosity, surface tension, or diffusion into the equations we evolve. These improvements (as well as including a physically-motivated equation of state) would help us validate our simulations of these and similar terrestrial experiments, but would be of little use in the astrophysical case, where viscosity and surface tension play no dynamic role.

Even with better diagnostics these types of experiments can only further validate the code if we model additional physics that are not relevant to the astrophysical problems of interest. As we mentioned above, flows within stars are expected to have a Reynolds number greater than 10^9 , for which the viscosity-free Euler equations are a better approximation, but such flows are impossible to achieve either in the laboratory or in current simulations. Therefore, we conclude that the terrestrial experiments we have simulated have served to build confidence in the hydrodynamics module in FLASH, but that there are certainly still limits on the strength of statements we can make about the validity of FLASH simulations of laboratory experiments. Our efforts at validation, however, despite presenting many challenges and leading to new questions, such as the effect of small scale structure on bulk properties of flows, have increased our confidence in the simulations produced by FLASH.

This work is supported in part by the U.S. Department of Energy under Grant No. B341495 to the Center for Astrophysical Thermonuclear Flashes at the University of Chicago, in part under the auspices of the U.S. Department of Energy by the University of California,

Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48, and in part by other U.S. Department of energy grants. K. Olson acknowledges partial support from NASA grant NAS5-28524, and L. J. Dursi acknowledges support by the Krell Institute CSGF. The work of T. Plewa was partly supported by the grant 2.P03D.014.19 from the Polish Committee for Scientific Research. Support for M. Zingale was provided by the Scientific Discovery through Advanced Computing (SciDAC) program of the DOE, grant number DE-FC02-01ER41176. The authors thank Mike Papka and the Argonne National Laboratory for visualization support. Finally, the authors thank Dale Pace for an insightful review of this manuscript.

REFERENCES

- AIAA 1998, Guide for the Verification and Validation of Computational Fluid Dynamics Simulations (AIAA Report G-077-1998; Reston, VA: American Institute of Aeronautics and Astronautics)
- Arnett, D., Fryxell, B., & Müller, E. 1989, *ApJ*, 341, L63
- Boehly, T. R., et al. 1995, *Rev. Sci. Instrum.* 66, 508
- Bradley, D. K., et al. 1998, *Phys. Plasmas* 5, 1870
- Calder, A. C., et al. 2000, in *Proc. Supercomputing 2000*, (IEEE Computer Soc.) <http://sc2000.org/proceedings/>
- Calder, A. C., et al. 2002, *ApJS*, in press
- Chandrasekhar, S. 1981, *Hydrodynamic and Hydromagnetic Stability* (New York: Dover)
- Colella, P., & Glaz, H. M. 1985, *J. Comput. Phys.*, 59, 264
- Colella, P., & Woodward, P. 1984, *J. Comput. Phys.*, 54, 174
- Dimonte, G., et al. 2002, in prep.
- Dimonte, G., & Schneider, M. 2000, *Phys. Fluids A*, 12, 304
- Dimonte, G., Morrison, J., Hulse, S., Nelson, D., Weaver, S., Susoeff, A., Hawke, R., Schneider, M., Batteaux, J., Lee, D., & Ticehurst, J. 1996, *Rev. Sci. Instrum.*, 67, 302
- Dimotakis, P. E. 2000, *J. Fluid Mech.*, 409, 69

- Duff, R. E., Harlow, F. H., & Hirt, C. W. 1962, *Phys. Fluids*, 4, 417
- Dursi, L. J., et al. 2002, In prep.
- Dyakov, S. P. 1954, *ZhETF* 27, 288 in Russian.
- Faber, T. E. 1995, *Fluid Dynamics for Physicists* (Cambridge: Cambridge Univ. Press), 293
- Freeman, N. C. 1955, *Proc. Roy. Soc. A*, 228, 341
- Fryxell, B. A., Müller, E., & Arnett, D. 1989, *Hydrodynamics and Nuclear Burning* (MPI Astrophys. Rep. 449; Garching: MPI Astrophys.)
- Fryxell, B., et al. 2000, *ApJS*, 131, 273
- Gehrz, R. D., Truran, J. W., Williams, R. E., & Starrfield, S. 1998, *PASP*, 110, 3
- Godunov, S. K. 1959, *Mat. Sbornik*, 47, 271
- Godunov, S. K., Zabrodin, A. V., & Prokopov, G. P. 1961, *U.S.S.R. Computational Math. and Math. Phys.*, 1, 1187
- Gropp, W., Lusk, E., & Skjellum, A. 1999, *Using MPI: Portable Parallel Programming with the Message Passing Interface*, 2nd ed. (Cambridge: MIT Press)
- Hillebrandt, W., & Niemeyer, J. C. 2000, *ARA&A*, 38, 191
- Holmes, R. L., Dimonte, G., Fryxell, B., Gittings, M. L., Grove, J. W., Schneider, M. S., Sharp, D. H., Velikovich, A. L., Weaver, R. P., & Zhang, Q. 1999, *J. Fluid. Mech.*, 389, 55
- Kane, J. O., Robey, H. F., Remington, B. A., Drake, R. P., Knauer, J., Ryutov, D. D., Louis, H., Teyssier, R., Hurricane, O., Arnett, D., Rosner, R., & Calder, A. 2001, *Phys. Rev. E*, 63, 055401 (R)
- Kifonidis, K., Plewa, T., Janka, H.-Th., & Müller, E. 2000, *ApJ*, 531, L123
- Khokhlov, A. M. 2001, *ApJ*, submitted (astro-ph/0008463)
- Khokhlov, A. M., Oran, E. S., & Wheeler, J. C. 1997, *ApJ*, 478, 678
- Lamb, D. Q. 2000, *ApJS*, 127, 395
- Landau, L. D., & Lifshitz, E. M. 1987, *Fluid Mechanics*, 2nd ed. (Oxford: Pergammon Press)

- Landen, O. L., et al. 2001, *Rev. Sci. Instrum.*, 72, 627
- Larson, J. T., & Lane, S. M. 1994, *J. Quant. Spect. Rad. Trans.* 51, 179
- Lewin, W. H. G., van Paradijs, J., & Taam, R. E. 1993, *Space Sci. Rev.*, 62, 223
- Lindl, J. D. 1998 *Inertial confinement Fusion* (Springer-Verlag: New York), 54
- Livio, M. 1994, *Mem. Soc. Astron. Ital.*, 65, 49
- MacNeice, P., Olson, K. M., Mobarrry, C., de Fainchtein, R., & Packer, C. 1999, NASA Technical Report CR-1999-209483
- MacNeice, P., Olson, K. M., Mobarrry, C., de Fainchtein, R., & Packer, C. 2000, *Comput. Phys. Commun.*, 126, 330
- Majda, A. 1984, *Compressible Fluid Flow and Systems of Conservation Laws in Several Space Variables* (New York: Springer-Verlag), 7
- Menikoff, R., & Plohr, B. J. 1989, *Rev. Mod. Phys.* 61, 75
- Meshkov, E. E. 1969, *Izv. Acad. Sci. USSR Fluid Dyn.* 4, 101
- Niemeyer, J. C. 1995, Ph.D. dissertation, Technical University of Munich.
- Niemeyer, J. C., & Hillebrandt, W. 1995, *ApJ*, 452, 769
- Nomoto, K., Yamaoka, H., & Shiegeyama, T. 1994, in *Supernovae, Les Houches Session LIV*, ed. S. Bludman, R. Mochkovitch, & J. Zinn-Justin (Amsterdam: Elsevier)
- Oberkampf, W. L. 1998, *Bibliography for Verification and Validation in Computational Simulation* (Sandia Report SAND98-2041; Albuquerque: Sandia National Laboratories)
- Oran, E. S., & Boris, J. P. 2001, *Numerical Simulation of Reactive Flow*, 2nd ed., (Cambridge: Cambridge University Press)
- Pilch, et al. 2001, *Guidlines for Sandia ASCI Verification and Validation Plans – Content and Format: Version 2.0* (Sandia Report SAND2000-3101; Albuquerque: Sandia National Laboratories)
- Richtmyer, R. D. 1960, *Comm. Pure Appl. Math.* 13, 297
- Roache, P. J. 1998, *Fundamentals of Computational Fluid Dynamics* (Albuquerque: Hermosa)

- Roache, P. J. 1998, *Verification and Validation in Computational Science and Engineering* (Albuquerque: Hermosa)
- Robey, H. F., et al. 2001, *Phys. Plasmas* 8, 2446
- Robey, H. F., et al. 2002, *Phys. Plasmas*, submitted
- Rosner, R., Alexakis, A., Young, Y.-N., Truran, J. W., & Hillebrandt, W. 2001, *ApJ*, 562, L177
- Ryutov, D., Drake, R. P., Kane, J. Liang, E., Remington, B. A., & Wood-Vasey, W. M. 1999, *ApJ*, 518, 821
- Schneider, M., Dimonte, G., & Remington, B. 1998, *Phys. Rev. Lett.*, 80, 3507
- Shara, M. M. 1989, *PASP*, 101, 5
- Shatz, H. et al. 2001, *Phys. Rev. Lett.*, 86, 3471
- Shu, C.-W., 1998 in "Advanced Numerical Approximation of Nonlinear Hyperbolic Equations", ed. A. Quarteroni (Springer:Berlin), pp. 325-432
- Soures, J. M., et al. 1996, *Phys. Plasmas*, 3, 2108
- Starrfield, S. 1989, in *Classical Novae*, ed. N. Evans & M. Bode, (New York: Wiley), 39
- Strang, G. 1968, *SIAM J. Num. Anal.*, 5, 506
- Taam, R. E. 1985, *Annu. Rev. Nucl. Part. Sci.*, 35, 1
- Taam, R. E., Woosley, S. E., Weaver, T. A., & Lamb, D. Q. 1993, *ApJ*, 413, 324
- Taylor, G. 1950, *Proc. Roy. Soc. Lond.*, A 201, 192
- www.top500.org
- Truran, J. W. 1982, in *Essays in Nuclear Astrophysics*, ed. C. A. Barnes, D. D. Clayton, & D. N. Schramm (Cambridge: Cambridge University Press), 467
- Turner, M. S. 2001, *PASP*, 113, 653
- Venn, J. 1880, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 9, 1
- Woodward, P. & Colella, P. 1984, *J. Comput. Phys.*, 54, 115

- Woosley, S. E., & Weaver, T. A. 1986, ARA&A, 24, 205
- Yee, H. C., Vinokur, M., & Djomehri, M. J. 2000, J. Comput. Phys., 162, 33
- Youngs, D. L. 1989, Physica D 37, 270
- Youngs, D. L. 1994, Lasers and Particle Beams, 12, 725
- Zingale, et al. 2002a, ApJS, in press